
Rule DAS611: Relatively large CI size was used for direct processing

Finding: CPExpert noticed that a relatively large CI size was used for direct processing. This finding applies only if SMF Type 42 (Data Set Statistics)¹ and SMF Type 64 (VSAM Statistics) records are available in a MXG performance data base.

Impact: This can have a LOW IMPACT, MEDIUM IMPACT, or HIGH IMPACT on the performance of applications referencing the VSAM data. The level of impact depends on the number of direct I/O operations.

Discussion: A VSAM file structure consists of one or more *Control Intervals (CIs)* and one or more *Control Areas (CAs)*.

- A **Control Interval** is a continuous area of direct access storage that VSAM uses to store logical records. When a logical record is read from direct access storage, the entire Control Interval containing the record is read into a VSAM buffer in virtual storage. The desired logical record is then transferred from the VSAM buffer to a user-defined buffer or work area. While logical records within a Control Interval may vary in length, all Control Intervals in a specific VSAM data set are of the same length.
- A **Control Area** contains one or more Control Intervals. The Control Intervals are grouped together into fixed-length contiguous areas of direct access storage. A VSAM data set is composed of one or more control areas.

Each VSAM data set is defined as a cluster of one or more components.

- The *data component* is the part of a VSAM data set, alternate index, or catalog that contains the data records. The minimum size of a Control Area for a data component is one track, and the maximum size is one cylinder of DASD storage.
- The *index component* is a collection of logically sequenced keys. A key is a value taken from a fixed defined field in each logical record in the VSAM data set. The key identifies the record's position in the data set. Using the index, VSAM is able to randomly retrieve a record from the data component when a request is made for a record with a certain key. The size of the Control Area for an index component is one track of DASD storage.

¹%LET TYPE42DS = Y; must be specified in USOURCE(GENGUIDE) must be specified in USOURCE(GENGUIDE) or USOURCE(DASGUIDE) to advise CPExpert that TYPE42DS is available.

The size of Control Intervals can vary from one VSAM data set to another, but all the Control Intervals of a particular data set component must be the same length. Control interval size affects record processing speed and storage requirements:

- Data sets with large control interval sizes require more buffer space in virtual storage.
- Data sets with large control interval sizes require fewer I/O operations to bring a given number of records into virtual storage, because fewer index records must be read.
- Free space is used more efficiently as control interval size increases relative to data record size. This is because there are fewer Control Interval splits and less wasted space.

The type of processing that is used should guide the choice of control interval size, particularly for the data component:

- **Sequential processing.** When sequential processing accounts for most of the accesses, a large Control Interval for the data component would normally be a good choice. This is because multiple records can be read into buffers and processed sequentially. For example, given a 16KB data buffer space, it is better to read two 8 KB Control Intervals with one I/O operation, than to read four 4 KB Control Intervals with two I/O operations.
- **Direct processing.** When direct processing accounts for most of the accesses, a small Control Interval for the data component is preferable. This is because only one logical record is retrieved at a time with direct access. Since a Control Interval normally contains several logical records, a large Control Interval would create unnecessary I/O overhead reading the logical records that would not be accessed. IBM suggests that a 4096 byte Control Interval normally would be appropriate for direct access.
- **Mixed processing.** If the processing is a mixture of sequential and direct, a small Control Interval for the data component with multiple buffers for sequential processing can be a good choice. The small Control Interval would reduce the I/O processing when using direct access to reference a logical record. The multiple buffers would allow read-ahead (and I/O efficiency) when using sequential processing.

After applying the screening criteria specified for VSAM data sets, and extracting SMF Type 64 information for those VSAM data sets, CPExpert examines SMF Type 42 (Data Set Statistics) information for the selected VSAM data sets. CPExpert uses the TYPE42DS information to compute

the percent of accesses to the data component of the VSAM data set that were direct, using the following algorithm:

$$\text{Percent direct accesses} = \frac{S42AMDRB}{S42AMSRB + S42AMDRB}$$

where: S42AMSRB = Blocks read using sequential access

S42AMDRB = Blocks read using direct access

CPEXpert produces Rule DAS611 when the percent of direct accesses for the data component was greater than the **PCTDIR** guidance variable in USOURCE(DASGUIDE), and the Control Interval size was greater than 4096 bytes. Additionally, CPEXpert verifies that the maximum logical records (maximum LRECL) is less than 50% of the Control Interval size. This verification is done to make sure that Rule DAS611 is not produced for VSAM data sets that have spanned records.

The default value for the PCTDIR guidance variable is 80%, so CPEXpert will produce Rule DAS611 when more than 80% of the accesses were direct for the data component, and the Control Interval size was more than 4096 bytes.

The following example illustrates the output from Rule DAS611:

RULE DAS611: RELATIVELY LARGE CI SIZE USED FOR DIRECT PROCESSING

VOLSER: CICS0E. A relatively large data Control Interval (CI) size was used for the VSAM data sets listed below, yet more than 80% of the accesses were for direct processing. You normally should define a CI size of 4KB for VSAM data sets used mostly for direct processing. The I/O RATE is for the time the data set was open.

SMF TIME STAMP	JOB NAME	VSAM DATA SET	..	I/O	-ACCESS TYPE (PCT)-	CI	
			..	RATE	SEQUENTIAL	DIRECT	SIZE
17:00,14OCT1997	DFHCOMDS	CICS410.PET.DFHCS.D.DATA.....		174.7	0.0	100.0	18432

While not shown in the above example, CPEXpert also shows the OPEN time for the VSAM data set. This normally is the duration of the current OPEN. If %LET VSAMSMRY=Y; was specified in USOURCE(DASGUIDE), the OPEN time represents the sum of the times the VSAM data was OPEN for all TYPE64 records in the performance data base.

Suggestion: If CPEXpert produces Rule DAS611, you should consider the following alternatives:

-
- Large Control Interval sizes for direct access normally degrade performance because the entire Control Interval (consisting of multiple logical records) is brought into storage with each access to a record. With direct processing, normally there is a low probability that more than one record in the Control Interval will be referenced. Consequently, you should consider decreasing the Control Interval size.
 - If the above action is not appropriate, you can change the **PCTDIR** guidance variable in USOURCE(DASGUIDE). Section 3 describes how to change the PCTSEQ guidance variable if you feel that Rule DAS611 is produced prematurely.
 - Alternatively, you can exclude the reported VSAM data sets from analysis. Section 3 describes how to exclude VSAM data sets from analysis. However, you should be aware that no analysis of potential VSAM problems will be performed on data sets that are excluded from analysis.

Reference: *DFSMS: Using Data Sets* (SC26-7339 for OS/390; SC26-7410 for z/OS)
Section 2.5.1.2: Data Control Interval Size